arXiv:1511.02626v2 [math.NA] 10 May 2016

$\mathcal{H}$-MATRIX BASED SECOND MOMENT ANALYSIS FOR ROUGH RANDOM FIELDS AND FINITE ELEMENT DISCRETIZATIONS

J. DÖLZ, H. HARBRECHT, AND M. PETERS

Abstract. We consider the efficient solution of strongly elliptic partial differential equations with random load based on the finite element method. The solution's two-point correlation can efficiently be approximated by means of an $\mathcal{H}$-matrix, in particular if the correlation length is rather short or the correlation kernel is non-smooth. Since the inverses of the finite element matrices which correspond to the differential operator under consideration can likewise efficiently be approximated in the $\mathcal{H}$-matrix format, we can solve the correspondent $\mathcal{H}$-matrix equation in essentially linear time by using the $\mathcal{H}$-matrix arithmetic. Numerical experiments for three dimensional finite element discretizations for several correlation lengths and different smoothness are provided. They validate the presented method and demonstrate that the computation times do not increase for non-smooth or shortly correlated data.

1. Introduction

A lot of problems in science and engineering can be modeled in terms of strongly elliptic boundary value problems. While these problems are numerically well understood for input data which are given exactly, these input data are often only available up to a certain accuracy in practical applications, e.g. due to measurement errors or tolerances in manufacturing processes. In recent years, it has therefore become more and more important to take these inaccuracies in the input data into account and model them as random input parameters.

The Monte Carlo approach, see e.g. [11] and the references therein, provides a straightforward approach to deal with these random data, but it has a relatively slow convergence rate which is only in the sense of the root mean square error. This, in turn, means that a large amount of samples has to be generated to obtain computational results with an acceptable accuracy, whereas the results still have a small probability of being too far away from the true solution. Therefore, in the past several years there have been presented multiple deterministic approaches to overcome this obstacle. For instance, random loads have been considered in [45, 48], random coefficients in [1, 2, 12, 17, 18, 38, 40, 42], and random domains in [33, 50].

For a domain $D \subset \mathbb{R}^d$ and a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, we consider the Dirichlet problem

$$\left.\begin{aligned} \mathcal{L}u(\omega,\mathbf{x}) &= f(\omega,\mathbf{x}) \quad && \text{for } \mathbf{x} \in D \\ u(\omega,\mathbf{x}) &= 0 && \text{for } \mathbf{x} \in \Gamma := \partial D \end{aligned}\right\} \quad \mathbb{P}\text{-almost surely} \tag{1}$$

with random load $f(\omega,\mathbf{x})$ and a strongly elliptic partial differential operator $\mathcal{L}$ of second order.

We can compute the solution's mean

$$\mathbb{E}_u(\mathbf{x}) := \int_\Omega u(\omega,\mathbf{x}) \,\mathrm{d}\mathbb{P}(\omega)$$

and also its two-point correlation

$$\mathrm{Cor}_u(\mathbf{x},\mathbf{y}) := \int_\Omega u(\omega,\mathbf{x})u(\omega,\mathbf{y}) \,\mathrm{d}\mathbb{P}(\omega)$$

if the respective quantities of the input data are known. Namely, the mean $\mathbb{E}_u$ satisfies

$$\mathcal{L}\mathbb{E}_u = \mathbb{E}_f \text{ in } D \quad \text{and} \quad \mathbb{E}_u = 0 \text{ on } \Gamma \tag{2}$$

2010 *Mathematics Subject Classification.* 60H35, 65C30, 65N30.

This research has been supported by the Swiss National Science Foundation (SNSF) through the project "$\mathcal{H}$-matrix based first and second moment analysis".

due to the linearity of the expectation and the differential operator $\mathcal{L}$. Taking into account the multi-linearity of the tensor product, one verifies by tensorizing (1) that

$$\begin{aligned} (\mathcal{L}\otimes\mathcal{L})\operatorname{Cor}_u &= \operatorname{Cor}_f && \text{in } D\times D,\\ (\mathcal{L}\otimes\operatorname{Id})\operatorname{Cor}_u &= 0 && \text{on } D\times\Gamma,\\ (\operatorname{Id}\otimes\mathcal{L})\operatorname{Cor}_u &= 0 && \text{on } \Gamma\times D,\\ \operatorname{Cor}_u &= 0 && \text{on } \Gamma\times\Gamma. \end{aligned} \tag{3}$$

From Cor_u, we can compute the variance $\mathbb{V}_u$ of the solution due to

$$\mathbb{V}_u(\mathbf{x}) = \operatorname{Cor}_u(\mathbf{x},\mathbf{x}) - \mathbb{E}_u(\mathbf{x})^2.$$

If a low-rank factorization of Cor_f is available, (3) can easily be solved by standard finite element methods. The existence of an accurate low-rank approximation is directly related to the spectral decomposition of the associated Hilbert-Schmidt operator

$$(\mathcal{C}_f\psi)(\mathbf{x}) := \int_D \operatorname{Cor}_f(\mathbf{x},\mathbf{y})\psi(\mathbf{y})\,\mathrm{d}\mathbf{y}. \tag{4}$$

Let $\operatorname{Cor}_f \in H^p(D\times D)$, then, according to [23], the eigenvalues of this Hilbert-Schmidt operator decay like

$$\lambda_m \lesssim m^{-2p/d} \text{ as } m\to\infty. \tag{5}$$

Unfortunately, the constant in this estimate involves the derivatives of Cor_f, or more precisely its $H^p(D)\otimes L^2(D)$-norm. The following consideration shows that this can lead to large constants in the decay estimate if the correlation length is small. Let the correlation kernel $k(r)$ depend only on the distance $r=\|\mathbf{x}-\mathbf{y}\|$. Then, the derivatives $\partial_{\mathbf{x}}^{\boldsymbol{\alpha}}\operatorname{Cor}_f(\mathbf{x},\mathbf{y})$ and $\partial_{\mathbf{y}}^{\boldsymbol{\alpha}}\operatorname{Cor}_f(\mathbf{x},\mathbf{y})$ of the correlation

$$\operatorname{Cor}_f(\mathbf{x},\mathbf{y}) = k\biggl(\frac{\|\mathbf{x}-\mathbf{y}\|}{\ell}\biggr),$$

involve the factor $\ell^{-|\boldsymbol{\alpha}|}$, leading to a constant ℓ^{-p} in the decay estimate of the eigenvalues (5). Thus, for small correlation length ℓ, a low-rank approximation of Cor_f becomes prohibitively expensive to compute.

Different approaches to tackle the solution of (3) have been considered in the following articles, where most of them have in common that they are in some sense based on a sparse tensor product discretization of the solution. For example, the computation of the second moment, i.e. Cor_u, has been considered for elliptic diffusion problems with random loads in [45] by means of a sparse tensor product finite element method. A sparse tensor product wavelet boundary element method has been used in [33] to compute the solution's second moment for elliptic potential problems on random domains. In [29, 32], the computation of the second moment was done by multilevel finite element frames. Recently, this concept has been simplified by using the combination technique, cf. [31]. Unfortunately, the sparse grid discretization needs to resolve the concentrated measure for short correlation lengths. This means that the number of hierarchies of the involved finite element spaces has to be doubled if the correlation length is halved to get the same accuracy.

The present article shall discuss a different approach to approximate the full tensor product discretization without losing its resolution properties. In [13], it has been demonstrated that the $\mathcal{H}$-matrix technique is a powerful tool to cope with Dirichlet data of low Sobolev smoothness if the problem is formulated as a boundary value problem. There, the similar behavior of two-point correlation kernels and boundary integral operators has been exploited. In [14], $\mathcal{H}$-matrix compressibility of the solution was proven also in case of local operators on domains. In the present article, we will combine this theoretical foundation with the $\mathcal{H}$-matrix technique used in [13] to efficiently solve (3) by the finite element method for a right hand side Cor_f with small correlation length or low Sobolev smoothness.

The general concept of $\mathcal{H}$-matrices and the corresponding arithmetic have at first been introduced in [24, 26]. $\mathcal{H}$-matrices are feasible for the data-sparse representation of (block-) matrices which can be approximated block-wise with low-rank and have originally been employed for the fast treatment of boundary integral equations as they arise in the boundary element method.

The rest of this article is organized as follows. In Section 2, we introduce the Galerkin discretization of the problem under consideration. Section 3 discusses the compressibility of discretized correlation kernels and the efficient solution of general correlation equations. Section 4 elaborates some specialties in the context of finite elements. Especially, we restate a phenomenon, called "weak admissibility", which produces a more data-sparse representation of the correlation matrices. In Section 5, we present numerical examples to validate and quantify the proposed method. Finally, in Section 6, we draw our conclusions from the theoretical findings and the numerical results.

2. Preliminaries

For the remainder of this article, let $D \subset \mathbb{R}^d$ be a bounded domain with Lipschitz boundary, $(\Omega, \mathcal{F}, \mathbb{P})$ a separable, complete probability space and $\mathcal{L}$ the linear differential operator of second order

$$(\mathcal{L}u)(\mathbf{x}) := -\operatorname{div}\big(\mathbf{A}(\mathbf{x}) \cdot \nabla u(\mathbf{x})\big) + c(\mathbf{x})u(\mathbf{x}). \tag{6}$$

The differential operator shall be strongly elliptic in the sense that

$$\underline{\alpha}\|\boldsymbol{\xi}\|_2^2 \leq \boldsymbol{\xi}^{\mathsf{T}}\mathbf{A}(\mathbf{x})\boldsymbol{\xi} \leq \overline{\alpha}\|\boldsymbol{\xi}\|_2^2 \text{ for all } \boldsymbol{\xi} \in \mathbb{R}^d$$

with coefficients $\mathbf{A} \in W^{1,\infty}\big(D, \mathbb{R}^{d\times d}\big)$ and $c \in L^\infty(D, \mathbb{R})$.

Under these assumptions, for a given load $f \in L^2_{\mathbb{P}}\big(\Omega, H^{-1}(D)\big)$, the Dirichlet problems

$$\begin{aligned} \mathcal{L}u(\omega, \mathbf{x}) &= f(\omega, \mathbf{x}) && \text{for } \mathbf{x} \in D \\ u(\omega, \mathbf{x}) &= 0 && \text{for } \mathbf{x} \in \Gamma := \partial D, \end{aligned}$$

are known to have a unique solution $u(\omega, \cdot) \in H^1_0(D)$ for $\mathbb{P}$-almost every $\omega \in \Omega$, cf. e.g. [19]. As a result, the mean $\mathbb{E}_u \in H^1_0(D)$ and the correlation $\operatorname{Cor}_u \in H^1_0(D) \otimes H^1_0(D)$ are uniquely determined.

Although the main focus of this article shall be on homogeneous boundary data, the existence and uniqueness results can straightforwardly be generalized to non-homogeneous boundary data $g \in L^2_{\mathbb{P}}\big(\Omega, H^{1/2}(\Gamma)\big)$, cf. also [19]. In this case, the tensorized problem (3) becomes

$$\begin{aligned} (\mathcal{L} \otimes \mathcal{L})\operatorname{Cor}_u &= \operatorname{Cor}_f && \text{in } D \times D, \\ (\mathcal{L} \otimes \operatorname{Id})\operatorname{Cor}_u &= \operatorname{Cor}_{f,g} && \text{on } D \times \Gamma, \\ (\operatorname{Id} \otimes \mathcal{L})\operatorname{Cor}_u &= \operatorname{Cor}_{g,f} && \text{on } \Gamma \times D, \\ \operatorname{Cor}_u &= \operatorname{Cor}_g && \text{on } \Gamma \times \Gamma, \end{aligned} \tag{7}$$

with the cross correlation

$$\operatorname{Cor}_{f,g}(\mathbf{x}, \mathbf{y}) = \int_\Omega f(\omega, \mathbf{x}) g(\omega, \mathbf{y}) \, \mathrm{d}\mathbb{P}(\omega) = \operatorname{Cor}_{g,f}(\mathbf{y}, \mathbf{x}), \quad \mathbf{x} \in D, \ \mathbf{y} \in \Gamma.$$

2.1. Galerkin discretization.

For the efficient numerical solution of (3), we use a finite element Galerkin scheme. To that end, we introduce a finite element space $V_N = \operatorname{span}\{\varphi_1, \dots, \varphi_N\} \subset H^1_0(D)$. It is assumed that the mesh which underlies this finite element space is quasi-uniform. The basis functions $\{\varphi_i\}_i$ are supposed to be locally and isotropically supported so that $\operatorname{diam}(\operatorname{supp}\varphi_i) \sim N^{-1/d}$. In particular, we can assign to each degree of freedom $i \in \{1, \dots, N\}$ a suitable point $\mathbf{x}_i \in D$, e.g. the barycenter of the support of the corresponding basis function or the corresponding Lagrangian interpolation point if nodal finite element shape functions are used.

The variational formulation of (3) is given as follows:

$$\begin{aligned} &\text{Find } \operatorname{Cor}_u \in H^1_0(D) \otimes H^1_0(D) \text{ such that} \\ &\big((\mathcal{L} \otimes \mathcal{L})\operatorname{Cor}_u, v\big)_{L^2(D\times D)} = (\operatorname{Cor}_f, v)_{L^2(D\times D)} \text{ for all } v \in H^1_0(D) \otimes H^1_0(D). \end{aligned}$$

By replacing the energy space $H^1_0(D) \otimes H^1_0(D)$ in this variational formulation by the finite dimensional ansatz space $V_N \otimes V_N$, we arrive at

$$\begin{aligned} &\text{Find } \operatorname{Cor}_{u,N} \in V_N \otimes V_N \text{ such that} \\ &\big((\mathcal{L} \otimes \mathcal{L})\operatorname{Cor}_{u,N}, v\big)_{L^2(D\times D)} = (\operatorname{Cor}_f, v)_{L^2(D\times D)} \text{ for all } v \in V_N \otimes V_N. \end{aligned} \tag{8}$$

A basis in $V_N \otimes V_N$ is formed by the set of tensor product basis functions $\{\varphi_i \otimes \varphi_j\}_{i,j}$. Hence, representing $\mathrm{Cor}_{u,N}$ by its basis expansion, yields

$$\mathrm{Cor}_{u,N} = \sum_{\ell,\ell'=1}^{N} u_{\ell,\ell'}(\varphi_\ell \otimes \varphi_{\ell'}).$$

Setting $\mathbf{C}_u := [u_{\ell,\ell'}]_{\ell,\ell'}$, we end up with the linear system of equations

$$(\mathbf{A} \otimes \mathbf{A})\,\mathrm{vec}(\mathbf{C}_u) = \mathrm{vec}(\mathbf{C}_f), \tag{9}$$

where $\mathbf{C}_f := \big[(\mathrm{Cor}_f, \varphi_\ell \otimes \varphi_{\ell'})_{L^2(D\times D)}\big]_{\ell,\ell'}$ is the discretized two-point correlation of the Dirichlet data f and $\mathbf{A} := \big[(\mathcal{L}\varphi_{\ell'}, \varphi_\ell)_{L^2(D)}\big]_{\ell,\ell'}$ is the system matrix of the second order differential operator (6). In (9), the tensor product has, as usual in connection with matrices, to be understood as the Kronecker product. Furthermore, for a matrix $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_n] \in \mathbb{R}^{m\times n}$, the operation $\mathrm{vec}(\mathbf{B})$ is defined as

$$\mathrm{vec}([\mathbf{b}_1, \ldots, \mathbf{b}_n]) := \begin{bmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_n \end{bmatrix} \in \mathbb{R}^{mn}.$$

For matrices $\mathbf{B} \in \mathbb{R}^{k\times n}$, $\mathbf{C} \in \mathbb{R}^{\ell\times m}$ and $\mathbf{X} \in \mathbb{R}^{m\times n}$, there holds the relation

$$(\mathbf{B} \otimes \mathbf{C})\,\mathrm{vec}(\mathbf{X}) = \mathrm{vec}(\mathbf{C}\mathbf{X}\mathbf{B}^\intercal).$$

Hence, we may rewrite (9) according to

$$\mathbf{A}\mathbf{C}_u\mathbf{A}^\intercal = \mathbf{C}_f. \tag{10}$$

2.2. **Non-homogeneous boundary conditions.** We shall briefly summarize on how to compute the solution of the non-homogeneous problem (7) and refer to [29] for a more detailed derivation. To that end, let us extend the finite element space V_N to $\overline{V}_N = V_N \oplus V_N^\Gamma \subset H^1(D)$, where V_N^Γ shall only contain basis functions with supports that intersect Γ.

To compute an approximate solution

$$\mathrm{Cor}_{u,N} = \sum_{\Theta,\Sigma\in\{V_N, V_N^\Gamma\}} \sum_{\varphi_i^\Theta\in\Theta} \sum_{\varphi_j^\Sigma\in\Sigma} u_{i,j}^{\Theta,\Sigma}\big(\varphi_i^\Theta \otimes \varphi_j^\Sigma\big) \in \overline{V}_N \otimes \overline{V}_N,$$

define the stiffness matrices

$$\mathbf{A}^{V_N,\Sigma} := \Big[\big(\mathcal{L}\varphi_j^\Sigma, \varphi_i^{V_N}\big)_{L^2(D)}\Big]_{\varphi_i^{V_N}\in V_N, \varphi_j^\Sigma\in\Sigma}, \quad \Sigma \in \big\{V_N, V_N^\Gamma\big\},$$

the mass matrix

$$\mathbf{G}^{V_N^\Gamma,V_N^\Gamma} := \big[(\varphi_j, \varphi_i)_{L^2(\Gamma)}\big]_{\varphi_i,\varphi_j\in V_N^\Gamma},$$

and the correlation matrices

$$\begin{aligned}
\mathbf{C}_f &:= \big[(\mathrm{Cor}_f, \varphi_\ell \otimes \varphi_{\ell'})_{L^2(D\times D)}\big]_{\varphi_\ell,\varphi_{\ell'}\in V_N}, \\
\mathbf{C}_{f,g} &:= \big[(\mathrm{Cor}_{f,g}, \varphi_\ell \otimes \varphi_{\ell'})_{L^2(D\times\Gamma)}\big]_{\varphi_\ell\in V_N,\varphi_{\ell'}\in V_N^\Gamma}, \\
\mathbf{C}_{g,f} &:= \big[(\mathrm{Cor}_{g,f}, \varphi_\ell \otimes \varphi_{\ell'})_{L^2(\Gamma\times D)}\big]_{\varphi_\ell\in V_N^\Gamma,\varphi_{\ell'}\in V_N}, \\
\mathbf{C}_g &:= \big[(\mathrm{Cor}_g, \varphi_\ell \otimes \varphi_{\ell'})_{L^2(\Gamma\times\Gamma)}\big]_{\varphi_\ell,\varphi_{\ell'}\in V_N^\Gamma}.
\end{aligned}$$

We observe that it holds $\mathbf{C}_{f,g} = \mathbf{C}_{g,f}^\intercal$. Moreover, we can represent the symmetric correlation matrix $\mathbf{C}_u$ of $\mathrm{Cor}_{u,N}$ by

$$\mathbf{C}_u = \begin{bmatrix} \mathbf{C}_u^{V_N,V_N} & \mathbf{C}_u^{V_N,V_N^\Gamma} \\ \mathbf{C}_u^{V_N^\Gamma,V_N} & \mathbf{C}_u^{V_N^\Gamma,V_N^\Gamma} \end{bmatrix}.$$

Thus, following [29], the solution $\mathbf{C}_u$ is given by the system

$$\begin{bmatrix} \mathbf{A}^{V_N,V_N} & \mathbf{A}^{V_N,V_N^\Gamma} \\ \mathbf{0} & \mathbf{G}^{V_N^\Gamma,V_N^\Gamma} \end{bmatrix} \begin{bmatrix} \mathbf{C}_u^{V_N,V_N} & \mathbf{C}_u^{V_N,V_N^\Gamma} \\ \mathbf{C}_u^{V_N^\Gamma,V_N} & \mathbf{C}_u^{V_N^\Gamma,V_N^\Gamma} \end{bmatrix} \begin{bmatrix} \mathbf{A}^{V_N,V_N} & \mathbf{A}^{V_N,V_N^\Gamma} \\ \mathbf{0} & \mathbf{G}^{V_N^\Gamma,V_N^\Gamma} \end{bmatrix}^\intercal = \begin{bmatrix} \mathbf{C}_f & \mathbf{C}_{f,g} \\ \mathbf{C}_{g,f} & \mathbf{C}_g \end{bmatrix},$$

which can successively be solved, see [29] for the details.

3. $\mathcal{H}$-matrix approximation of correlation kernels

The linear system (10) of equations has N^2 unknowns and is therefore not directly solvable if N is large due to memory and time consumption. Thus, an efficient compression scheme and a powerful arithmetic is needed to obtain its solution. In the following, we will restrict ourselves to *asymptotically smooth* correlation kernels Cor_f, i.e. correlation kernels satisfying the following definition.

Definition 3.1. *Let $k\colon \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$. The function k is called* asymptotically smooth *if for some constants $c, r > 0$ and $q \in \mathbb{R}$ holds*

$$\left|\partial_{\mathbf{x}}^{\boldsymbol{\alpha}} \partial_{\mathbf{y}}^{\boldsymbol{\beta}} k(\mathbf{x}, \mathbf{y})\right| \leq c \frac{(|\boldsymbol{\alpha}| + |\boldsymbol{\beta}|)!}{r^{|\boldsymbol{\alpha}|+|\boldsymbol{\beta}|}} \|\mathbf{x} - \mathbf{y}\|_2^{-2-2q-|\boldsymbol{\alpha}|-|\boldsymbol{\beta}|} \tag{11}$$

independently of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$.

Examples for asymptotically smooth correlation kernels are the Matérn kernels, which include especially the Gaussian kernel, cf. [39, 43] and the references therein. A main feature of such asymptotically smooth correlation kernels is that they exhibit a data-sparse representation by means of $\mathcal{H}$-matrices, cf. e.g. [5, 8, 25].

$\mathcal{H}$-matrices rely on local low-rank approximations of a given matrix $\mathbf{X} \in \mathbb{R}^{N \times N}$. For suitable index sets $\nu, \nu' \subset \{1, \dots, N\}$, a matrix block $\mathbf{X}|_{\nu \times \nu'}$ can be approximated by a rank-k matrix. This approximation can be represented in factorized form $\mathbf{X}|_{\nu \times \nu'} \approx \mathbf{Y}\mathbf{Z}^\intercal$ with factors $\mathbf{Y} \in \mathbb{R}^{\#\nu \times k}$ and $\mathbf{Z} \in \mathbb{R}^{\#\nu' \times k}$. Hence, if $k \ll \min\{\#\nu, \#\nu'\}$, the complexity for storing the block is considerably reduced. The construction of the index sets is based on the *cluster tree*.

3.1. Cluster tree.

A tree $\mathcal{T} = (V, E)$ with vertices V and edges E is called a *cluster tree* for the set $\{1, \dots, N\}$ if the following conditions hold.

- $\{1, \dots, N\}$ is the *root* of $\mathcal{T}$.
- All non-empty $\sigma \in V \setminus \mathcal{L}(\mathcal{T})$ are the disjoint union of their sons.
- The number of sons of each vertex $\sigma \in V \setminus \mathcal{L}(\mathcal{T})$ is at least two.

Hereby, we define the set of leaves of $\mathcal{T}$ by

$$\mathcal{L}(\mathcal{T}) := \{\sigma \in V : \sigma \text{ has no sons}\}.$$

We say $\sigma' \in V$ is a *son* of $\sigma \in V$ if $(\sigma, \sigma') \in E$. The *level* of $\sigma \in \mathcal{T}$ is its distance of the root, i.e. the number of son relations that are required for traveling from $\{1, \dots, N\}$ to σ. In the following, we will refer to the vertices as *clusters* and denote the k-th vertex on level j by $\sigma_{j,k}$.

The construction of the cluster tree is based on the *support* of the clusters. The support Υ_σ of a cluster σ is defined as the union of the supports of the basis functions corresponding to their elements, that is

$$\Upsilon_\sigma = \bigcup_{i \in \sigma} \Upsilon_i \text{ where } \Upsilon_i := \operatorname{supp} \varphi_i \text{ for all } i \in \{1, \dots, N\}.$$

For computing complexity bounds, the cluster tree should match the following additional requirements, uniformly as $N \to \infty$:

- The cluster tree is a balanced binary tree in the sense that the maximal level satisfies $J \sim \log_2 N$ and the number of sons of the clusters $\sigma_{j,k}$ is two whenever the cluster is not a leaf.
- The diameter of the support $\Upsilon_{\sigma_{j,k}}$ is local with respect to the level j, i.e. $\operatorname{diam} \Upsilon_{\sigma_{j,k}} \sim 2^{-j/d}$. Moreover, the number $\#\sigma_{j,k}$ of indices contained in a cluster $\sigma_{j,k}$ on level j scales approximately like 2^{J-j}, i.e. $\#\sigma_{j,k} \sim 2^{J-j}$.

The cluster tree $\mathcal{T}$ with the indicated terms should be given for our further considerations. A common algorithm for its construction is based on a hierarchical subdivision of the point set which is associated with the basis functions, cf. e.g. [5, 8, 25]. We begin by embedding the point set $\{x_1, \dots, x_N\}$ in a top-level cuboid. This cuboid is subsequently subdivided into two cuboids of the same size where the corresponding clusters are described by the points in each cuboid. This process is iterated until a cuboid encloses less than a predetermined number of points.

3.2. **$\mathcal{H}$-Matrix approximation.** $\mathcal{H}$-matrices have originally been invented in [24, 26] and are a generalization of cluster techniques for the rapid solution of boundary integral equations such as the fast multipole method [22], the mosaic skeleton approximation [47], or the adaptive cross approximation [3].

For the discretization of an asymptotically smooth correlation, we introduce a partition of its domain of definition which separates smooth and non-smooth areas of the kernel function. It is based on the following

Definition 3.2. *Two clusters σ and σ' are called* η-admissible *if*

$$\max\{\operatorname{diam}(\Upsilon_\sigma), \operatorname{diam}(\Upsilon_{\sigma'})\} \leq \eta \operatorname{dist}(\Upsilon_\sigma, \Upsilon_{\sigma'}) \tag{12}$$

holds for some fixed $\eta > 0$.

We can obtain the set of admissible blocks by means of a recursive algorithm: Starting with the root $(\sigma_{0,0}, \sigma_{0,0})$, the current cluster pair is checked for admissibility and, if admissible, added to the set $\mathcal{F}$ which corresponds to the the correlation kernels's *farfield*. Otherwise, the admissibility check will be performed on all possible pairs of son clusters of the two original clusters. When we arrive at a pair of inadmissible leaf clusters, it is added to the set $\mathcal{N}$ which corresponds to the correlation kernel's *nearfield*. The set $\mathcal{B} = \mathcal{F} \cup \mathcal{N}$ obviously inherits a tree a tree structure from the recursive construction of $\mathcal{F}$ and $\mathcal{N}$ and is called the *block cluster tree*, see [5, 8, 25].

With the definition of the block cluster tree at hand, we are finally in the position to introduce $\mathcal{H}$-matrices.

Definition 3.3. *The set $\mathcal{H}(\mathcal{B}, k)$ of* $\mathcal{H}$-matrices *of maximal block rank k is defined according to*

$$\mathcal{H}(\mathcal{B}, k) := \big\{\mathbf{X} \in \mathbb{R}^{N\times N} : \operatorname{rank}\big(\mathbf{X}|_{\sigma\times\sigma'}\big) \leq k \text{ for all } \big(\sigma, \sigma'\big) \in \mathcal{F}\big\}.$$

Note that all nearfield blocks $\mathbf{X}|_{\sigma\times\sigma'}$, $(\sigma, \sigma') \in \mathcal{N}$, are allowed to be full matrices.

In accordance with [5, 8, 25], the storage cost of an $\mathcal{H}$-matrix $\mathbf{X} \in \mathcal{H}(\mathcal{B}, k)$ is $\mathcal{O}(kN \log N)$. Here, for asymptotically smooth correlation kernels, the rank k depends logarithmically on the desired approximation accuracy ε, which in turn depends usually on the degrees of freedom N. These remarks pertain to the approximation of an *explicitly* given, asymptotically smooth correlation kernel k, such as Cor_f in (3).

The compressibility of an *implicitly* given correlation kernel, such as Cor_u in (3), has been studied in [14] for the case of smooth domains D. We restate the main theorem for the setting of the present article which employs that the Hilbert-Schmidt operator (4), related with the correlation kernel Cor_f, is in general a *pseudo-differential operator*, see e.g. [34, 35, 36, 46] and the references therein.

Theorem 3.4. *In the domain D with analytic boundary ∂D, assume that the correlation kernel Cor_f in (3) gives rise to an operator $\mathcal{C}_f \in OPS^{\theta}_{cl,s}$, i.e. to a classical pseudo-differential operator with symbol $a_f(\mathbf{x}, \boldsymbol{\xi})$ of order θ and of Gevrey class $s \geq 1$ in the sense of* [10, Def. 1.1]. *Assume further that the coefficients of the differential operator $\mathcal{L}$ are smooth. Then, the correlation kernel Cor_u of (3) is the Schwartz kernel of an operator $\mathcal{C}_u \in OPS^{\theta-4}_{cl,s}$.*

Moreover, the kernel $\operatorname{Cor}_u(\mathbf{x}, \mathbf{y})$ of the correlation operator $\mathcal{C}_u$ is smooth in $D \times D$ outside of the diagonal $\Delta := \{(\mathbf{x}, \mathbf{y}) \in D \times D : \mathbf{x} = \mathbf{y}\}$ and there holds the pointwise estimate

$$\begin{aligned} &\forall \boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{N}_0^n, (\mathbf{x}, \mathbf{y}) \in (D \times D)\backslash\Delta\colon \\ &\qquad |\partial_{\mathbf{x}}^{\boldsymbol{\alpha}} \partial_{\mathbf{y}}^{\boldsymbol{\beta}} \operatorname{Cor}_u(\mathbf{x}, \mathbf{y})| \leq c\mathscr{A}^{|\boldsymbol{\alpha}+\boldsymbol{\beta}|}(|\boldsymbol{\alpha}|!)^s \boldsymbol{\beta}! \|\mathbf{x} - \mathbf{y}\|_2^{-\theta-n-|\boldsymbol{\alpha}|-|\boldsymbol{\beta}|+4} \end{aligned} \tag{13}$$

for some constants c and $\mathscr{A}$ which depend only on D and on a_f.

Obviously, for $s = 1$, equality (13) directly implies condition (11) for the asymptotic smoothness of Cor_u, allowing us to approximate Cor_u by the means of $\mathcal{H}$-matrices. In particular, [14] provides also some numerical evidence that this result could likely be extended to Lipschitz domains.

An example of correlation kernels for Cor_f satisfying the condition of this theorem for $s \geq 1$ is the Matérn class of kernels. We refer to [14] for more details on how to verify the assumptions of the theorem for other correlation kernels.

3.3. **$\mathcal{H}$-Matrix arithmetic and iterative solution.** An important feature of $\mathcal{H}$-matrices is that efficient algorithms for approximate matrix arithmetic operations are available. For two $\mathcal{H}$-matrices $\mathbf{H}_1, \mathbf{H}_2 \in \mathcal{H}(\mathcal{B}, k)$, the approximate *matrix-matrix addition* $\mathbf{H}_1$+$\mathbf{H}_2 \in \mathcal{H}(\mathcal{B}, k)$ can be performed in $\mathcal{O}(k^2 N \log N)$ operations while the approximate *matrix-matrix multiplication* $\mathbf{H}_1$*$\mathbf{H}_2 \in \mathcal{H}(\mathcal{B}, k)$ can be performed in $\mathcal{O}(k^2 N \log^2 N)$ operations. Both of these operations are essentially block matrix algorithms with successive recompression schemes. Moreover, employing the recursive block structure, the approximate inversion or the approximate computation of the LU-decomposition within $\mathcal{H}(\mathcal{B}, k)$ can also be performed in only $\mathcal{O}(k^2 N \log^2 N)$ operations. We refer the reader to [5, 8, 21, 25, 26] for further results and implementation details.

In the context of correlation equations, this approximate $\mathcal{H}$-matrix arithmetic has successfully been used in [13] to solve a problem, similar to (3), which has been discretized by the boundary element method. Then, the matrix $\mathbf{A}$ in (10) corresponds to the stiffness matrix from the boundary element method which can naturally be approximated by the $\mathcal{H}$-matrix technique. The resulting matrix equation has been solved using an iterative solver based on *iterative refinement*, cf. [20, 41, 49], which we are also going to employ here. This method has originally been introduced in [49] for the improvement of solutions to linear systems of equations based on the LU-factorization.

Having all matrices in (10) represented by $\mathcal{H}$-matrices, the solution can then be approximated as follows. Let $\hat{\mathbf{A}}^{-1} \in \mathcal{H}(\mathcal{B}, k)$ be an approximate inverse to $\mathbf{A}$, e.g. computed from $\mathbf{A}$ by the $\mathcal{H}$-matrix arithmetic. Starting with the initial guess $\mathbf{C}_u^{(0)} = \hat{\mathbf{A}}^{-1}\mathbf{C}_f\hat{\mathbf{A}}^{-\intercal}$, we iterate

$$\boldsymbol{\Theta}^{(i)} = \mathbf{C}_f - \mathbf{A}\mathbf{C}_u^{(i)}\mathbf{A}^{\intercal}, \quad \mathbf{C}_u^{(i+1)} = \mathbf{C}_u^{(i)} + \hat{\mathbf{A}}^{-1}\boldsymbol{\Theta}^{(i)}\hat{\mathbf{A}}^{-\intercal}, \quad i = 0, 1, \dots. \tag{14}$$

The approximate inverse can either be computed explicitly or by employing an approximate LU-decomposition with forward and backward substitution algorithms. In both cases, the idea of the iterative refinement stays the same: The residual $\boldsymbol{\Theta}^{(i)}$ is computed with a higher precision than the correction $\hat{\mathbf{A}}^{-1}\boldsymbol{\Theta}^{(i)}\hat{\mathbf{A}}^{-\intercal}$. This yields an improved approximation to the solution in each step. Note that this algorithm also algebraically coincides with an undamped preconditioned Richardson iteration, see e.g. [44].

In the following, we will elaborate how this approach can be realized in the context of the finite element method.

4. $\mathcal{H}$-matrices in the context of finite elements

Although a finite element matrix has a sparse structure, its inverse is generally fully populated. Nevertheless, the inverse exhibits a data-sparse structure in the sense that it is $\mathcal{H}$-matrix compressible. We recall the main concepts from the literature, see e.g. [4, 7, 15, 16, 25].

4.1. **General concepts.** A rough argument for the $\mathcal{H}$-matrix compressibility of the inverse makes use of the Green's function G of $\mathcal{L}$. Let $\delta_{\mathbf{x}}$ denote the delta distribution at the point $\mathbf{x}$ and let $G\colon \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ satisfy

$$\mathcal{L}_{\mathbf{y}} G(\mathbf{x}, \mathbf{y}) = \delta_{\mathbf{x}} \text{ and } G(\cdot, \mathbf{y})|_{\Gamma} = 0.$$

Then, the solution of

$$\begin{aligned} \mathcal{L}u(\mathbf{x}) &= f(\mathbf{x}) && \text{for } \mathbf{x} \in D, \\ u(\mathbf{x}) &= 0 && \text{for } \mathbf{x} \in \Gamma, \end{aligned}$$

can be represented by

$$u(\mathbf{x}) = (\mathcal{L}^{-1} f)(\mathbf{x}) = \int_D G(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) \, \mathrm{d}\mathbf{y}, \quad \mathbf{x} \in D.$$

If the Green's function is analytic away from the diagonal, e.g. in the case of constant coefficients of $\mathcal{L}$, we can approximate the Green's function away from the diagonal by local expansions of the kind

$$G(\mathbf{x}, \mathbf{y}) \approx \sum_{i=1}^{k} a(\mathbf{x}) b(\mathbf{y}),$$

which is the theoretical basis for an $\mathcal{H}$-matrix approximation, see [5, 8, 25].

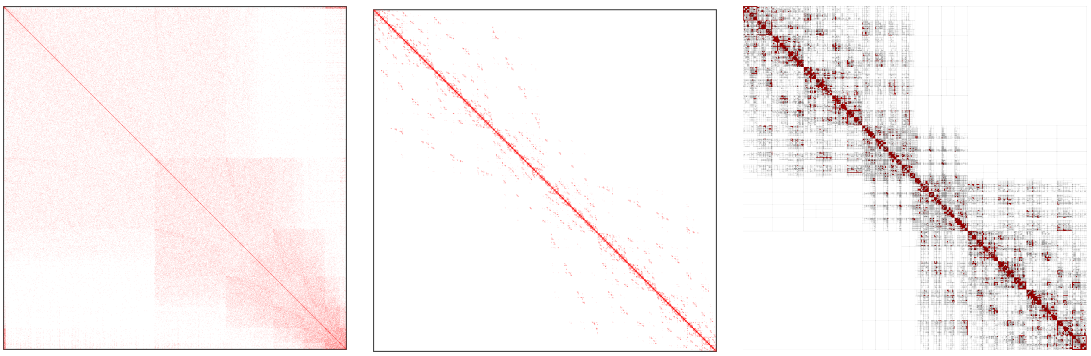

sparse FEM-matrix reordered FEM-matrix $\mathcal{H}$-matrix representation

FIGURE 1. Sparsity pattern of a 3D finite element matrix, its reordered finite element matrix, and the corresponding $\mathcal{H}$-matrix. Red blocks in the $\mathcal{H}$-matrix correspond to the nearfield, white blocks correspond to the empty farfield.

However, one of the advantages of the finite element method is that it can be applied also in case of non-constant coefficients. In [7], a proof was presented to guarantee the existence of an $\mathcal{H}$-matrix approximation to the inverse of the finite element stiffness matrix even in the case of L^∞-diffusion coefficients and the other coefficients set to zero. This result was then extended in [4, 6] to allow all coefficients to be L^∞, providing the theoretical foundation for an $\mathcal{H}$-matrix approximation to the inverse of the differential operator from (6). While these first results hold up to the finite element discretization error, the results have recently been improved in [15, 16] to hold without additional error.

It remains to explain how to actually compute an $\mathcal{H}$-matrix approximation to the inverse of a finite element stiffness matrix. To that end, note that a necessary condition for an entry $\mathbf{A}_{ij}$ in the finite element matrix to be non-zero is that $\Upsilon_i \cap \Upsilon_j \neq \emptyset$, i.e. the intersection of the corresponding supports of the basis functions is non-empty. This yields together with the η-adminissibility condition (12) that all entries of a finite element matrix have η-inadmissible supports, i.e. they are contained in the nearfield of an $\mathcal{H}$-matrix. A sparse finite element matrix can therefore be represented as an $\mathcal{H}$-matrix by reordering the index set corresponding to the clustering scheme introduced in Section 3.1 and inserting the non-zero entries into the nearfield. An illustration of this procedure can be found in Figure 1.

Having the finite element matrix represented by an $\mathcal{H}$-matrix, the approximate inverse can be computed by using the block algorithms of the $\mathcal{H}$-matrix arithmetic in $\mathcal{O}(k^2 N \log^2 N)$ operations, cf. [5, 8, 25]. However, since the approximate LU-decomposition is needed for the computation of the inverse, there is also another possibility where one uses forward and backward substitution algorithms to obtain a faster and more accurate algorithm to apply the inverse, cf. [5, 8, 25] and the references therein. We note especially that the two factors of the LU-decomposition are $\mathcal{H}$-matrix compressible, cf. [15, 16], and the overall complexity for this approach is still $\mathcal{O}(k^2 N \log^2 N)$, but with lower constants.

4.2. **Weak admissibility.** Approximate $\mathcal{H}$-matrix representations for the inverse or LU-factorizations of finite element matrices have been used to construct preconditioners for iterative solvers, see e.g. [5] and the references therein. In [27], it was observed that the computation of an approximate inverse can be considerably sped up by replacing the η-admissibility condition (12) by the following *weak admissibility condition.*

Definition 4.1. *Two clusters σ and σ' are called* weakly admissible *if $\sigma \neq \sigma'$.*

We observe immediately that an η-admissible block cluster is also weakly admissible. Thus, by replacing the η-admissibility condition by the weak admissibility, we obtain a much coarser partition of the $\mathcal{H}$-matrix. This leads to smaller constants in the storage and computational complexity, cf. [27]. Each row and each column of the finite element matrix has only $\mathcal{O}(1)$ entries.

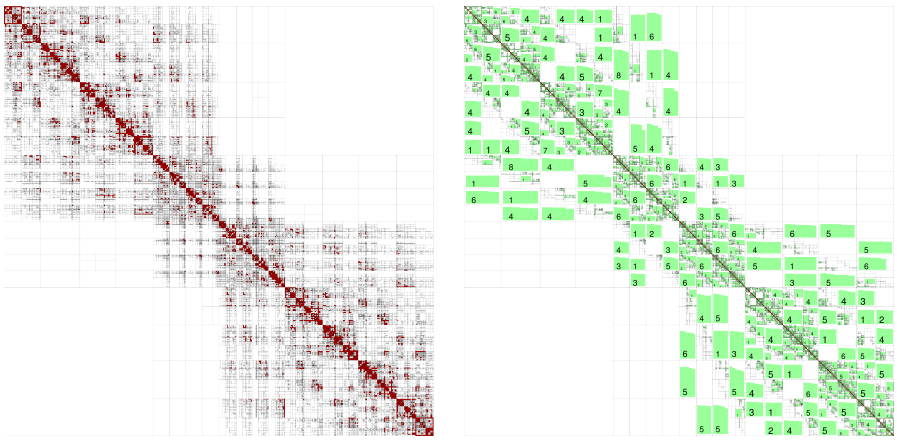

η-admissible FEM-matrix weakly admissible FEM-matrix

FIGURE 2. Comparison of the partition for η-admissibility and for weak admissibility. Red blocks correspond to full matrices, green blocks correspond to low-rank matrices with inscribed rank, and white blocks are zero.

Thus, inserting the finite element matrix into a weakly admissible $\mathcal{H}$-matrix structure, the off-diagonal blocks of the $\mathcal{H}$-matrix have low-rank.

By partitioning the matrix according to the weak admissibility condition, we cannot ensure the exponential convergence of fast black box low-rank approximation techniques as used for boundary element matrices. For example, the adaptive cross approximation (ACA) relies on an admissibility condition similar to (12) to ensure exponential convergence, cf. [3]. Instead, the authors of [27] suggest to assemble a weakly admissible matrix block according to the η-admissibility condition and transform it on-the-fly to a low-rank matrix to obtain a good approximation.

The behavior of the ranks of the low-rank matrices in weakly admissible partitions compared to η-admissible partitions is not fully understood yet. Suppose that k_η is an upper bound for the ranks corresponding to an η-admissible partition and suppose that k_w shall be an upper bound for the ranks to a weakly admissible partition. In [27], it is proven for one dimensional finite element discretizations that one should generally choose

$$k_w = Lk_\eta,$$

in order to obtain the same approximation accuracy in the weakly admissible case as in the η-admissible case. Here, L is a constant which depends on the depth of the block cluster tree and thus logarithmically on N. Already in the same article, the authors remark in the numerical examples that this bound on k_w seems to be too pessimistic and one could possibly choose

$$k_w = c_{\eta\to w} k_\eta, \tag{15}$$

where $c_{\eta\to w} \in [2, 3.5]$.

One can try to reduce the influence of the unknown behavior of the weak admissibility condition by mixing it with the η-admissibility. In the software package HLib, cf. [9], the authors use the η-admissibility for all block clusters with a block size larger than a given threshold and apply the weak admissibility condition for block clusters $\sigma \times \sigma'$ which are below that threshold provided that the condition $a_i^\sigma < \frac{a_i^{\sigma'}+b_i^{\sigma'}}{2} < b_i^\sigma$ or $a_i^{\sigma'} < \frac{a_i^\sigma+b_i^\sigma}{2} < b_i^{\sigma'}$ is satisfied for the corresponding bounding boxes $\prod_{i=1}^3 [a_i^\mu, b_i^\mu]$, $\mu = \sigma, \sigma'$, in at most one coordinate direction. The impact of this specific admissibility condition is illustrated in Figure 2.

5. NUMERICAL RESULTS

All the computations in the following experiments have been carried out on a single core of a computing server with two Intel(R) Xeon(R) E5-2670 CPUs with a clock rate of 2.60GHz and a main memory of 256GB. For the $\mathcal{H}$-matrix computations, we use the software package HLib, cf. [9],

and for the finite element discretization the Partial Differential Equation Toolbox of MATLAB which employs piecewise linear finite elements. The two libraries are coupled together in a single C-program, cf. [37], using the MATLAB Engine interface.

5.1. **Experimental setup.** To obtain computational efficiency and to keep the ranks of the low-rank matrices under control, HLib imposes an upper threshold k_η for the ranks in the case of an η-admissible $\mathcal{H}$-matrix and a lower threshold $n_{\min}$ for the minimal block size. For the application of the weak admissibility condition, we rely on the criterion of HLib, which considers the weak admissibility condition only if one of the index sets of the block cluster $\sigma\times\sigma'$ has a cardinality below 1,024 and the condition $a_i^\sigma < \frac{a_i^{\sigma'}+b_i^{\sigma'}}{2} < b_i^\sigma$ or $a_i^{\sigma'} < \frac{a_i^{\sigma}+b_i^{\sigma}}{2} < b_i^{\sigma'}$ is satisfied for the corresponding bounding boxes $\prod_{i=1}^3[a_i^\mu, b_i^\mu]$, $\mu=\sigma,\sigma'$, in at most one coordinate direction. Otherwise, the η-admissibility is used instead. In the case of a weakly admissible matrix block, HLib imposes an upper threshold of $k_w = 3k_\eta$, setting $c_{\eta\to w} = 3$ in (15). For our experiments, we choose $\eta = 2$, $k_\eta = 20$, $n_{\min} = 50$ and employ a geometric cluster strategy, i.e. the cluster strategy from the end of Section 3.1. For the application of the approximate inverse in the iterative refinement (14), we compute an approximate LU-decomposition and use forward and backward substition. The iterative refinement is stopped if the absolute error of the residual in the Frobenius norm is smaller 10^{-6}.

In the following examples, we want, besides other aspects, to study the influence of the weak admissibility condition for the partitioning of the different $\mathcal{H}$-matrices. Namely, we successively want to replace the η-admissibility by the weak admissibility as described in Table 1 in order to lower the constants hidden in the complexity of the $\mathcal{H}$-matrix arithmetic and thus to improve the computation time. For the discretization of the correlation kernel Cor_f, we will always use ACA.

Case	Operator	
	$\mathcal{L}$ and $\mathcal{L}^{-1}$	Cor_f and Cor_u
all-η	η-admissibility	η-admissibility
weak-FEM	weak admissibility	η-admissibility
weak-Cor	weak admissibility	weak admissibility

TABLE 1. The three combinations of the admissibility conditions used for the partition of the $\mathcal{H}$-matrices.

The *all-η* case is the canonical case and has also been investigated in case of the boundary element method in [13]. The *weak-FEM* case is a first relaxation to apply the weak admissibility condition. This is justified, since the stiffness matrix $\mathbf{A}$ can exactly be represented as a weakly admissible $\mathcal{H}$-matrix and the iterative refinement only involves an approximate LU-decomposition for the application of the approximate inverse. Hence, we expect at most an influence on the quality of the approximate LU-decomposition and thus on the number of iterations in the iterative refinement. We have therefore to investigate if the additional iterations are compensated by the faster $\mathcal{H}$-matrix arithmetic.

The aforementioned cases have in common that they rely on the asymptotic smoothness of Cor_f and Cor_u and the η-admissibility which leads to exponential convergence of the $\mathcal{H}$-matrix approximation. In the case *weak-Cor*, we want to examine if there is some indication that the weak admissibility could possibly also be considered for the partition of the $\mathcal{H}$-matrices for Cor_f and Cor_u. To that end, we approximate Cor_f with ACA relative to the η-admissibility partition and convert it on-the-fly to the partition of the weak admissibility, as proposed in [27].

The following numerical examples are divided in two parts. In the first part, we demonstrate the convergence of the presented method by comparing it to a low-rank reference solution computed with the pivoted Cholesky decomposition, cf. [30]. In the second part, we will demonstrate that the presented method works also well in the case of correlation kernels with low Sobolev smoothness or small correlation length, where no low-rank approximations exist and sparse grid approximations fail to resolve the correlation length. Note that in both examples, due to the non-locality of the

correlation kernels and the Green's function, the computed system matrices are smaller than usual for the finite element method. In particular, the unknown in the system of equations (10) is also a matrix.

5.2. **Tests for the iterative solver.** Due to the recompression schemes in the block matrix algorithms of the $\mathcal{H}$-matrix arithmetic, it is not directly clear if the presented solver provides convergence. Still, it can be shown that some iterative $\mathcal{H}$-matrix schemes converge up to a certain accuracy, cf. [28]. In the following, we want to demonstrate for a specific example that our iterative scheme provides indeed convergence.

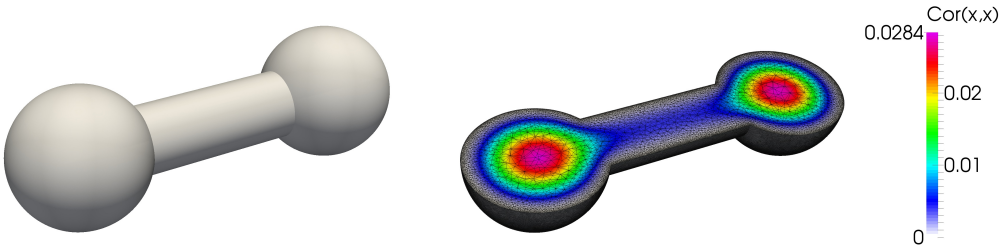

FIGURE 3. The dumbbell geometry and its meshed cross section with the reference solution for the Matérn-5/2 kernel.

On the dumbbell geometry pictured in Figure 3, we consider $\mathcal{L} = -\Delta$ in (3) and the Matérn-5/2 kernel as input correlation Cor_f, i.e. for $r = \|\mathbf{x} - \mathbf{y}\|_2$, we set

$$\mathrm{Cor}_f(\mathbf{x}, \mathbf{y}) = \left(1 + \sqrt{5}\frac{r}{\ell} + \frac{5}{3}\frac{r^2}{\ell^2}\right) \exp\left(-\sqrt{5}\frac{r}{\ell}\right),$$

where $\ell \approx \mathrm{diam}(D)$ denotes the correlation length. The conversion of the finite element matrix to an $\mathcal{H}$-matrix for the dumbbell geometry has already been illustrated in Figure 1. Whereas, the difference between the η-admissibility and the weak admissibility is illustrated in Figure 2.

Level	1	2	3	4	5	6	Reference mesh
Mesh points	95	351	1,438	6,539	28,659	120,947	472,745
DoF	11	105	611	3,673	18,178	81,782	330,615

TABLE 2. Mesh points and degree of freedoms (DoF) for different levels on the dumbbell geometry.

For determining a reference solution, we compute a low-rank approximation $\mathbf{C}_f \approx \mathbf{L}_f\mathbf{L}_f^\intercal$ with the pivoted Cholesky decomposition as proposed in [30]. The numerical solution $\mathbf{C}_u$ of (10) is then given by

$$\mathbf{C}_u \approx \mathbf{L}_u\mathbf{L}_u^\intercal,$$

where $\mathbf{L}_u$ solves $\mathbf{A}\mathbf{L}_u = \mathbf{L}_f$. To compute the error of the $\mathcal{H}$-matrix solution, we compare the correlation's trace $\mathrm{Cor}_{u,N}|_{\mathbf{x}=\mathbf{y}}$ of the $\mathcal{H}$-matrix approximation to the trace $\mathrm{Cor}_{u,N}|_{\mathbf{x}=\mathbf{y}}$ derived from the pivoted Cholesky decomposition on a finer reference mesh. We refer to Table 2 for more details on the meshes under consideration.

The appropriate norm for error measurements is the $W^{1,1}$-norm, for which we can expect a convergence rate of the mesh size h, stemming from the convergence rate of the finite element method in the H^1-norm. Figure 4 shows that we indeed reach this rate for all three cases of admissibility which are considered in Table 1. In fact, the observed errors are equal in the first few digits.

We are also interested in the quality of the approximate LU-decomposition $\mathbf{A} \approx \hat{\mathbf{L}}\hat{\mathbf{U}}$. We use a built-in function of the HLib to estimate the deviation $\hat{\mathbf{L}}\hat{\mathbf{U}} - \mathbf{A}$ in the spectral norm by ten power iterations, which is a good indicator for the approximation quality of the LU-decomposition of

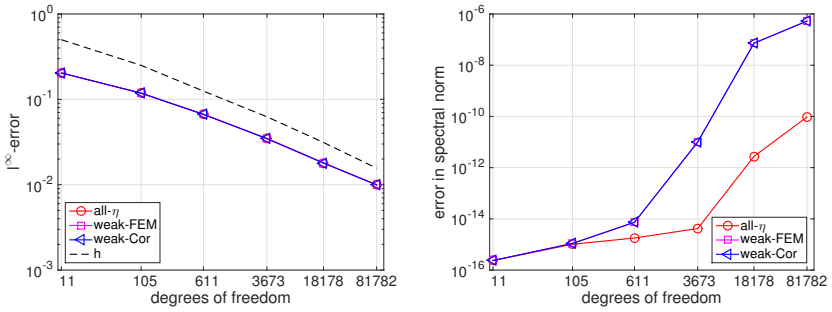

FIGURE 4. $W^{1,1}$-error for $\mathrm{Cor}_{u,N}|_{\mathbf{x}=\mathbf{y}}$ on the dumbbell (left) and the deviation $\hat{\mathbf{L}}\hat{\mathbf{U}} - \mathbf{A}$ in the spectral norm for fixed rank (right).

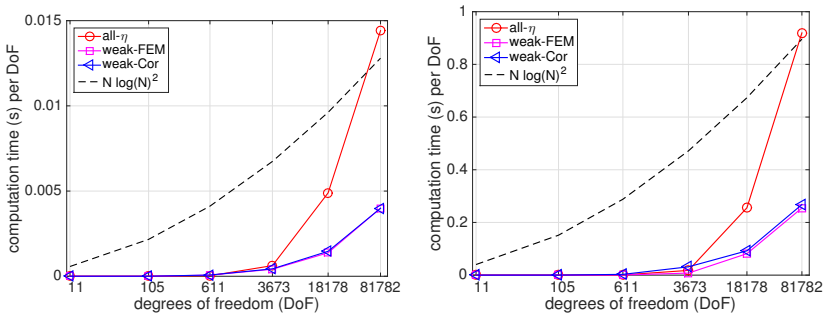

FIGURE 5. Computation times in seconds for the approximate LU-decomposition of the system matrix $\mathbf{A}$ (left) and for the iterative refinement (right) on the dumbbell geometry.

the finite element matrix. The estimated errors are plotted in Figure 4. Note that the observed behavior is in contrast to the behavior typically observed for preconditioning, cf. e.g. [5], since we do not increase the rank with the number of unknowns. We can see that the LU-decomposition is most accurate in the all-η case. Still, only one iteration is needed in the iterative refinement in all cases. When it comes to computation times, Figure 5 and Tables 3 and 4 illustrate that all cases of admissibility under consideration seem to run in essentially linear complexity, where the weak admissibility condition leads to a small speed-up. Figure 6 illustrates the required average and maximal ranks needed for the computations, whereas Figure 7 illustrates the amount of storage needed per degree of freedom. For reasons of performance, HLib allocates the worst case scenario for the ranks, thus in the latter case only weak and η-admissibility for a single $\mathcal{H}$-matrix have to be considered.

Level	1	2	3	4	5	6
full-η	2.9e-05	0.000366	0.019645	2.29641	88.9833	1178.77
weak-FEM	2.4e-05	0.000267	0.041767	1.51114	25.0998	326.771
weak-Cor	1.8e-05	0.000287	0.044564	1.60242	26.6771	326.177

TABLE 3. Computation times in seconds to compute the approximate LU-decomposition of the finite element matrix on the dumbbell geometry.

Having verified the convergence of our solver, we now want to consider different correlation lengths and smoothness.

Level	1	2	3	4	5	6
full-η	0.000105	0.00966	0.468599	65.914	4646.4	75196.7
weak-FEM	8.6e-05	0.007681	0.342041	28.4528	1480.83	20839.3
weak-Cor	8.3e-05	0.011927	1.78032	115.977	1679.27	21907.4

TABLE 4. Computation times in seconds for the iterative refinement on the dumbbell geometry.

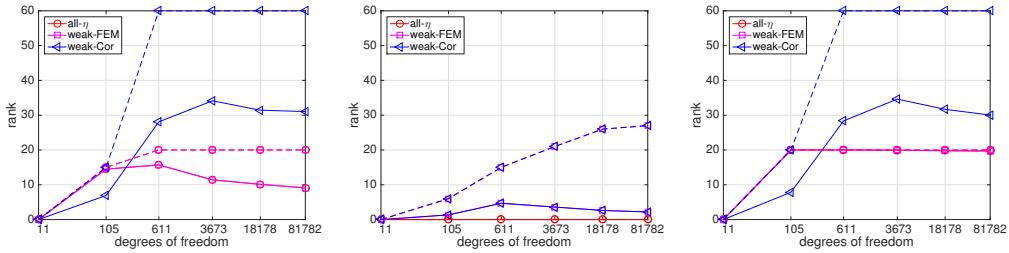

FIGURE 6. Required ranks for the prescribed correlation (left), the LU-decomposition (middle) and the solution correlation (right). The straight lines indicate the average ranks, whereas the dashed lines illustrate the maximal rank attained.

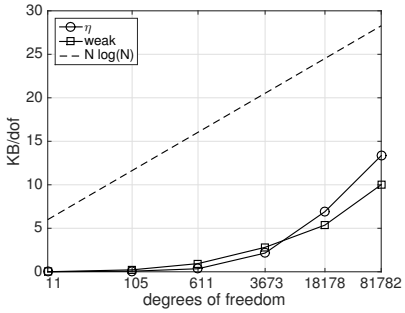

FIGURE 7. Allocated storage for the η-admissibity and the weak admissibility per degree of freedom. The allocated storage is independent of the content of the matrix.

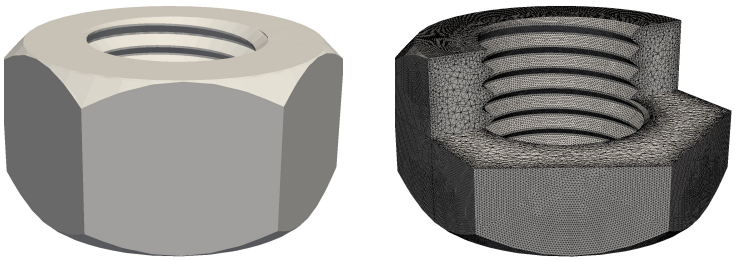

FIGURE 8. The screw-nut geometry (left) and its meshed cross section (right).

5.3. **Rough correlations.** In the second part of the numerical experiments, we employ correlation kernels with smaller correlation lengths and lower regularity such that low-rank approximations would become prohibitively expensive and sparse grid approaches would fail to resolve the concentrated measure.

We consider the screw-nut geometry pictured in Figure 8 which is discretized by a mesh with 240,877 vertices, 162,738 degrees of freedom, and an average element diameter of $h/\operatorname{diam}(D) \approx$

$\ell/\operatorname{diam}(D)$	1	0.1	0.01	0.001
Exponential kernel	70,321.3	70,691.8	64,551.5	58,868.8
Gaussian kernel	67,198.8	70,233.1	62,570.8	55,836.3

TABLE 5. Computation times in seconds for the iterative refinement on the screw-nut geometry for the exponential kernel and the Gaussian kernel with different correlation lengths.

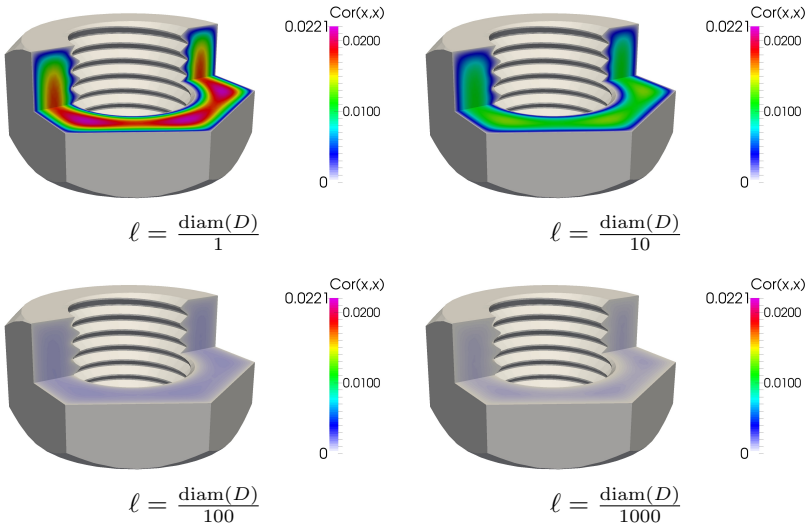

FIGURE 9. Cross sections through the screw-nut geometry for the exponential kernel with different correlation lengths ℓ.

0.0226. We choose $\mathcal{L} = -\Delta$ in (3) and either the Gaussian kernel as input correlation Cor_f, i.e.

$$\operatorname{Cor}_f(\mathbf{x},\mathbf{y}) = \exp\left(-\frac{\|\mathbf{x}-\mathbf{y}\|^2}{2\ell^2}\right),$$

or the exponential kernel, i.e.

$$\operatorname{Cor}_f(\mathbf{x},\mathbf{y}) = \exp\left(-\frac{\|\mathbf{x}-\mathbf{y}\|}{\ell}\right).$$

Herein, $\ell > 0$ denotes the correlation length.

In the following, we want to demonstrate that the presented method is very well suited for small correlation lengths ℓ. We therefore choose the correlation lengths

$$\ell \in \left\{\frac{\operatorname{diam}(D)}{1}, \frac{\operatorname{diam}(D)}{10}, \frac{\operatorname{diam}(D)}{100}, \frac{\operatorname{diam}(D)}{1,000}\right\}$$

for both, the Gaussian kernel and the exponential kernel, and compute the corresponding correlation of the solution Cor_u of (3).

The computation time for the approximate LU-decomposition is around 1,000 seconds, whereas the computation times for the iterative refinement are contained in Table 5. Note that the iterative refinement needed only one iteration. Moreover, it is clearly visible that the computation times do not grow with decreasing ℓ. In fact, it seems that the method works even better for small correlation lengths as the correlation matrix $\mathbf{C}_f$ tends towards a diagonal matrix and ranks in the off-diagonal blocks decrease. The smoothness of the prescribed Cor_f does not seem to have an impact on the computation time.

The cross sections found in Figure 9 illustrate the different behaviour of the correlation's trace $\operatorname{Cor}_u|_{\mathbf{x}=\mathbf{y}}$ for the different correlation lengths in case of the exponential kernel. The related results for the Gaussian kernel are presented in Figure 10.

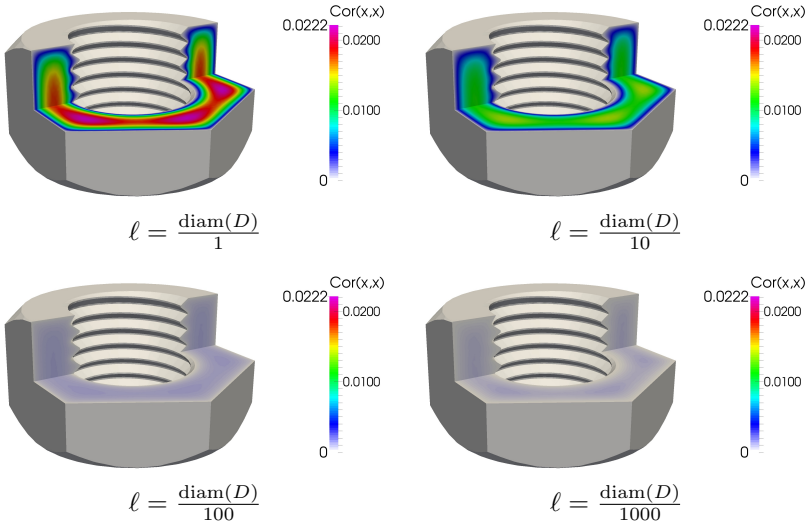

FIGURE 10. Cross sections through the screw-nut geometry for the Gaussian kernel with different correlation lengths ℓ.

6. Conclusion

We considered the solution of strongly elliptic partial differential equations with random load by means of the finite element method. Approximating the full tensor approach by means of $\mathcal{H}$-matrices, we employed the $\mathcal{H}$-matrix technique to efficiently discretize the non-local correlation kernel of the data and approximate the LU-decomposition of the finite element stiffness matrix. The corresponding $\mathcal{H}$-matrix equation has then been efficiently solved in essentially linear complexity by the $\mathcal{H}$-matrix arithmetic.

Compared to sparse grid or low-rank approximations, the proposed method does not suffer in case of roughly correlated data from large constants in the complexity estimates or the lack of resolution of the roughness. This was shown by numerical experiments on a non-trivial three dimensional geometry. Indeed, neither the computation times nor the storage requirements do increase for correlation kernels with short correlation length. It was moreover demonstrated that the use of the weak admissibility for the partition of the $\mathcal{H}$-matrix improves the constants in the computational complexity without having a significant impact to the solution accuracy.

References

1. I. Babuška, F. Nobile, and R. Tempone, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM Journal on Numerical Analysis **45** (2007), no. 3, 1005–1034.
2. I. Babuška, R. Tempone, and G. Zouraris, *Galerkin finite element approximations of stochastic elliptic partial differential equations*, SIAM Journal on Numerical Analysis **42** (2004), no. 2, 800–825.
3. M. Bebendorf, *Approximation of boundary element matrices*, Numerische Mathematik **86** (2000), no. 4, 565–589.
4. ———, *Efficient inversion of the Galerkin matrix of general second-order elliptic operators with nonsmooth coefficients*, Mathematics of Computation **74** (2005), no. 251, 1179–1199.
5. ———, *Hierarchical matrices*, Lecture Notes in Computational Science and Engineering, vol. 63, Springer Berlin Heidelberg, 2008 (English).
6. ———, *Low-rank approximation of elliptic boundary value problems with high-contrast coefficients*, SIAM Journal on Mathematical Analysis **48** (2016), no. 2, 932–949.
7. M. Bebendorf and W. Hackbusch, *Existence of $\mathcal{H}$-matrix approximants to the inverse FE-matrix of elliptic operators with L^∞-coefficients*, Numerische Mathematik **95** (2003), no. 1, 1–28 (English).
8. S. Börm, *Efficient Numerical Methods for Non-local Operators*, EMS Tracts in Mathematics, vol. 14, European Mathematical Society (EMS), Zürich, 2010. MR 2767920 (2012i:65001)
9. S. Börm, L. Grasedyck, and W. Hackbusch, *Hierarchical matrices*, Tech. Report 21, Max Planck Institute for Mathematics in the Sciences, 2003.

10. L. Boutet de Monvel and P. Krée, *Pseudo-differential operators and Gevrey classes*, Annales de l'Institut Fourier (Grenoble) **17** (1967), no. fasc. 1, 295–323. MR 0226170 (37 #1760)
11. R. Caflisch, *Monte Carlo and quasi-Monte Carlo methods*, Acta Numerica **7** (1998), 1–49.
12. M. Deb, I. Babuška, and J. Oden, *Solution of stochastic partial differential equations using Galerkin finite element techniques*, Computer Methods in Applied Mechanics and Engineering **190** (2001), no. 48, 6359–6372. MR 1870425 (2003g:65009)
13. J. Dölz, H. Harbrecht, and M. Peters, *$\mathcal{H}$-matrix accelerated second moment analysis for potentials with rough correlation*, Journal of Scientific Computing **65** (2015), no. 1, 387–410 (English).
14. J. Dölz, H. Harbrecht, and C. Schwab, *Covariance regularity and $\mathcal{H}$-matrix approximation for rough random fields*, Preprint 2014-11, Mathematisches Institut, Universität Basel (2014).
15. M. Faustmann, *Approximation inverser Finite Elemente- und Randelementematrizen mittels hierarchischer matrizen*, PHD Thesis, Technische Universität Wien, 2015.
16. Markus Faustmann, Jens Markus Melenk, and Dirk Praetorius, *$\mathcal{H}$-matrix approximability of the inverses of FEM matrices*, Numerische Mathematik **131** (2015), no. 4, 615–642.
17. P. Frauenfelder, C. Schwab, and R.A. Todor, *Finite elements for elliptic problems with stochastic coefficients*, Computer Methods in Applied Mechanics and Engineering **194** (2005), no. 2-5, 205–228.
18. R. Ghanem and P. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer, New York, 1991.
19. D. Gilbarg and N.S. Trudinger, *Elliptic partial differential equations of second order*, Springer, Berlin-Heidelberg, 1977.
20. G.H. Golub and C.F. Van Loan, *Matrix Computations*, 4th ed., Johns Hopkins University Press, Baltimore, 2012.
21. L. Grasedyck and W. Hackbusch, *Construction and arithmetics of $\mathcal{H}$-matrices*, Computing **70** (2003), no. 4, 295–334.
22. L. Greengard and V. Rokhlin, *A fast algorithm for particle simulations*, Journal of Computational Physics **73** (1987), no. 2, 325–348.
23. M. Griebel and H. Harbrecht, *Approximation of bi-variate functions: singular value decomposition versus sparse grids*, IMA Journal of Numerical Analysis **34** (2014), no. 1, 28–54.
24. W. Hackbusch, *A sparse matrix arithmetic based on $\mathcal{H}$-matrices part I: Introduction to $\mathcal{H}$-matrices*, Computing **62** (1999), no. 2, 89–108.
25. ———, *Hierarchical Matrices: Algorithms and Analysis*, Springer, Heidelberg, 2015.
26. W. Hackbusch and B.N. Khoromskij, *A sparse $\mathcal{H}$-matrix arithmetic. General complexity estimates*, Journal of Computational and Applied Mathematics **125** (2000), no. 12, 479–501.
27. W. Hackbusch, B.N. Khoromskij, and R. Kriemann, *Hierarchical matrices based on a weak admissibility criterion*, Computing **73** (2004), no. 3, 207–243 (English).
28. W. Hackbusch, B.N. Khoromskij, and E.E. Tyrtyshnikov, *Approximate iterations for structured matrices*, Numerische Mathematik **109** (2008), no. 3, 365–383.
29. H. Harbrecht, *A finite element method for elliptic problems with stochastic input data*, Applied Numerical Mathematics **60** (2010), no. 3, 227–244.
30. H. Harbrecht, M. Peters, and R. Schneider, *On the low-rank approximation by the pivoted Cholesky decomposition*, Applied Numerical Mathematics **62** (2012), 28–440.
31. H. Harbrecht, M. Peters, and M. Siebenmorgen, *Combination technique based k-th moment analysis of elliptic problems with random diffusion*, Journal of Computational Physics **252** (2013), 128–141.
32. H. Harbrecht, R. Schneider, and C. Schwab, *Multilevel frames for sparse tensor product spaces*, Numerische Mathematik **110** (2008), no. 2, 199–220.
33. ———, *Sparse second moment analysis for elliptic problems in stochastic domains*, Numerische Mathematik **109** (2008), no. 3, 385–414 (English).
34. L. Hörmander, *The Analysis of Linear Partial Differential Operators. I*, Classics in Mathematics, Springer, Berlin, 2003, Distribution theory and Fourier analysis, Reprint of the second (1990) edition. MR 1996773
35. ———, *The Analysis of Linear Partial Differential Operators. III*, Classics in Mathematics, Springer, Berlin, 2007, Pseudo-differential operators, Reprint of the 1994 edition.
36. G.C. Hsiao and W.L. Wendland, *Boundary Integral Equations*, Applied Mathematical Sciences, vol. 164, Springer, Berlin, 2008. MR 2441884 (2009i:45001)
37. B.W. Kernighan and D.M. Ritchie, *The C Programming Language*, Prentice-Hall, Upper Saddle River, NJ, 1988.
38. B.N. Khoromskij and C. Schwab, *Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs*, SIAM Journal on Scientific Computing **33** (2011), no. 1, 364–385.
39. B. Matérn, *Spatial variation*, Springer Lecture Notes in Statistics, Springer, New York, 1986.
40. H. Matthies and A. Keese, *Galerkin methods for linear and nonlinear elliptic stochastic partial differential equations*, Computer Methods in Applied Mechanics and Engineering **194** (2005), no. 12-16, 1295–1331.
41. C.B. Moler, *Iterative refinement in floating point*, Journal of the ACM **14** (1967), no. 2, 316–321.
42. F. Nobile, R. Tempone, and C.G. Webster, *An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM Journal on Numerical Analysis **46** (2008), no. 5, 2411–2442. MR 2421041 (2009c:65331)

43. C.E. Rasmussen and C.K.I. Williams, *Gaussian processes for machine learning (adaptive computation and machine learning)*, The MIT Press, Cambridge, 2005.
44. Y. Saad, *Iterative methods for sparse linear systems*, second ed., Society for Industrial and Applied Mathematics, 2003.
45. C. Schwab and R.A. Todor, *Sparse finite elements for elliptic problems with stochastic loading*, Numerische Mathematik **95** (2003), no. 4, 707–734.
46. M.E. Taylor, *Pseudodifferential Operators*, Princeton Mathematical Series, vol. 34, Princeton University Press, Princeton, NJ, 1981. MR 618463 (82i:35172)
47. E. Tyrtyshnikov, *Mosaic sceleton approximation*, Calcolo **33** (1996), 47–57.
48. T. von Petersdorff and C. Schwab, *Sparse finite element methods for operator equations with stochastic data*, Applications of Mathematics **51** (2006), no. 2, 145–180.
49. J.H. Wilkinson, *Rounding errors in algebraic processes*, Prentice Hall, Englewood Cliffs, NJ, 1963.
50. D. Xiu and D.M. Tartakovsky, *Numerical methods for differential equations in random domains*, SIAM Journal on Scientific Computing **28** (2006), no. 3, 1167–1185.

Departement Mathematik und Informatik, Fachbereich Mathematik, Universität Basel, Spiegelgasse 1, 4051 Basel, Switzerland

E-mail address: {Juergen.Doelz, Helmut.Harbrecht, Michael.Peters}@unibas.ch